WORKING PAPERS

Collective Commitment

Christian Roessler
Sandro Shelegia
Bruno Strulovici

June 2015

Working Paper No: 1507

DEPARTMENT OF ECONOMICS

UNIVERSITY OF VIENNA

Collective Commitment*

Christian Roessler Sandro Shelegia Bruno Strulovici

June 28, 2015

Abstract

We consider collective decisions made by agents whose preferences and power depend on past events and decisions. Faced with an inefficient equilibrium and an opportunity to commit to a policy, can the agents reach an agreement on such a policy? Under an intuitive condition linking power structures in the dynamic setting and at the commitment stage, the answer is negative: when the condition holds, the only agreement that may be reached at the outset, if any, coincides with the equilibrium without commitment. The condition is also necessary: when it fails, as in the case of a single time-inconsistent agent, commitment is valuable for some payoffs. We apply our result to explain inefficient collective decisions in the contexts of investment in a public good, hiring, and reform. *JEL: D70, H41, C70*

1 Introduction

Consider a group of persons making successive decisions as their preferences, information, and political power evolve over time. Given a chance to commit to a policy at the outset, can this group agree on a policy that improves the equilibrium outcome? How does the answer depend on the voting rules used in the dynamic setting, and on the rules used to select a commitment?

The importance of these questions is highlighted by the frequent inefficiency of political outcomes, such as the inability to implement needed reforms (Fernandez and Rodrik (1991), the use of short-sighted monetary or fiscal policies (Kydland and Prescott (1977) and Battaglini and Coate (2008)), the stability of unpopular regimes (Acemoglu and Robinson (2005)), and the invocation of slippery slope arguments (Volokh (2003)). In these contexts, it is legitimate to ask whether commitment could or even should be used to address such

*Roessler: California State University, East Bay (e-mail: christian.roessler@csueastbay.edu); Shelegia: University of Vienna (e-mail: sandro.shelegia@univie.ac.at); Strulovici: Northwestern University (e-mail: b-strulovici@northwestern.edu). We are grateful to Wiola Dziuda, Georgy Egorov, Jeff Ely, Daniel Garcia, Michael Greinecker, Karl Schlag, Stephen Schmidt, Joel Watson and numerous seminar participants for their comments. Strulovici gratefully acknowledges financial support from an NSF CAREER Award (Grant No. 1151410) and a fellowship form the Alfred P. Sloan Foundation.

inefficiency. Wouldn't the potential benefits be large? Couldn't constitutions, laws, and other contracts guarantee the validity of these commitments?

To illustrate the issue, consider a city in need of developing some infrastructure, but whose support for the construction is subject to shifts in public opinion and political power. In Rome, for example, there is a wide agreement that the metro system is underdeveloped.[1] However, new developments have been stalled, and Roman officials have explained the difficulties as follows: building a new metro line may result in the discovery of antique ruins which many citizens would find too valuable to destroy, resulting in metro construction to be abandoned.[2] This potential withdrawal of support may deter those citizens who value the metro but not additional ruins from supporting the costly metro construction in the first place. The deadlock could seemingly be resolved by a citywide commitment to complete metro construction regardless of what is found underground.[3] As it turns out, however, such a commitment is majority-preferred to the status quo *if and only if* it is itself majority dominated by the policy consisting of starting the metro line and then abandoning it if valuable ruins are uncovered during construction. This policy is, in turn, dominated by the status quo, generating a Condorcet cycle.

The situation is depicted on Figure 1. Voters are equally divided in three groups (A, B, C) whose terminal payoffs are indicated on the figure. The population first decides whether to dig a metro line (*Y*es or *N*o). If they do, an antiquity may be discovered with probability q, in which case society chooses between pursuing the *M*etro or preserving the antique *T*reasure. Decisions are made according to the simple majority rule (general rules are considered later). If the treasure were surely available ($q = 1$), restoring it would be the Condorcet winner among the three terminal outcomes and yield the highest utilitarian welfare. For $q \in (2/3, 1]$, metro construction followed by restoration in case the treasure is found (policy "YT") is the unique equilibrium, while constructing the metro unconditionally (YM) is the unique equilibrium for $q \in [0, 1/3)$, as we show in the Appendix. Policy YM also beats the status quo, N, in majority voting and yields a higher utilitarian welfare.

For $q \in (1/3, 2/3)$, however, the equilibrium policy is the status quo, N, because voters in group A deem too high the risk that ruins are found and the metro, abandoned, while voters in group C find the probability of ruins too low to justify digging the metro. Since N is majority dominated by policy YM, and YT gets a majority support (groups B and C) over YM, this yields a Condorcet cycle among policies: $YM < YT < N < YM$. Therefore, allowing the city to commit to a state-contingent plan at the outset is unlikely to resolve the

[1] The Roman underground has two lines, 49 stations, that serve a metropolitan area of 3.4 million residents. As a comparison, Berlin is similar in size but has 173 subway stations. Madrid, about one-and-a-half times as large as Rome, has 300. The only other European metropolis with such a low number of metro stations is Athens, which has 3.3 million inhabitants and 33 metro stops.

[2] The chairman of Roma Metropolitane SpA, Enrico Testa, was quoted saying: "There are treasures that are underground that would stay buried forever, but as soon as we uncover them, our work gets blocked." (Kahn (2007))

[3] One could preserve the most valuable pieces; in fact, the argument does not rely on the destruction of any ruins, as long as preservation is costly.

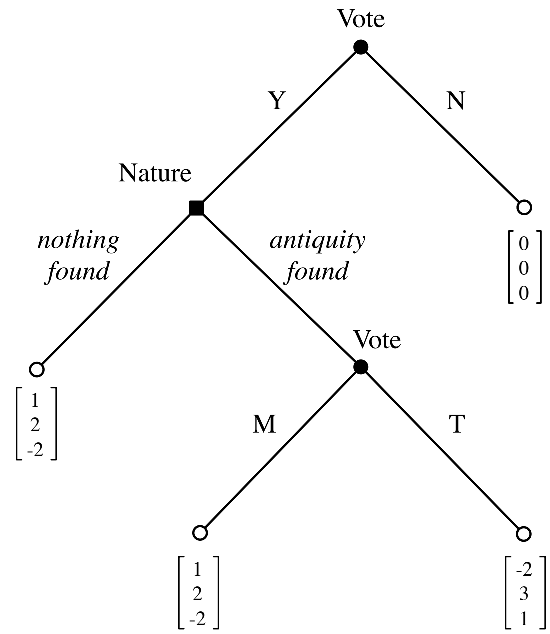

Figure 1: The Roman Metro. Circular nodes indicate majority-based decisions.

problem. It only leads to disagreements over the plan to follow and, in particular, will not rule out the status quo as a viable option.

This paper addresses the following questions: When can collective commitment be used to improve dynamic political equilibria? And how does the value of commitment depend on the voting rule used and, more generally, on the political power of each agent? One must distinguish agents' power on dynamic collective decisions absent commitment (dynamic stage) from their power on the ranking of state-contingent policies (commitment stage). The relationship between these two forms of power is instrumental in determining the value of commitment as well as who gets to benefit from it.

To develop an understanding of these questions, we first consider the case in which all decisions, in both the dynamic and the commitment stages, are made according to the simple majority rule. The main conclusion of the Roman metro example then holds in full generality: given any payoffs, either the equilibrium is undominated or there is a Condorcet cycle at the commitment stage which includes both the equilibrium and the policy dominating it. As a result, the dominating policy is no more valuable, according to the democratic criterion of majority voting, than the equilibrium itself.

Instituting a commitment stage thus replaces any problem of *inefficiency* by one of *indeterminacy*, even when commitments carry no administrative or other contractual costs and are perfectly credible. Commitment is thus never valuable because it is either unnecessary, when the equilibrium is socially preferred to all other state-contingent policies, or it is impossible to agree upon which commitment to choose.

We then consider general structures of political power. These may include the use of supermajority rules for some decisions and heterogenous allocations of power across agents. These details, it turns out, do not matter *per se* for the value of commitment. Instead, this value is determined by a *power consistency* condition relating political power at the dynamic and commitment stages. When power consistency holds, the introduction of commitment suffers from the same problem as in the case of majority voting: it is either unnecessary or leads to indeterminacy. Furthermore, the power consistency condition is not only sufficient but also necessary for this result: when it is violated, one may find a preference profile and a policy that dominates not just the equilibrium but also all other policies at the commitment stage.

We focus on dynamic settings where decisions are binary in each period and may be made according to arbitrary voting rules. The focus on binary decisions eliminates "local" Condorcet cycles in each period and thus also an important potential source of indeterminacy which may confuse the main points of the paper.[4]

Power consistency is then defined as the following requirement: Consider two policies which are identical except for the decision made in a given period and for a given state (or subset of states) in this period. Then, the social ranking between these two policies must be determined by the same set of winning coalitions as the one arising in the dynamic game when that decision is reached. The condition thus rules out situations in which a subset of persons could impose one policy over another at the commitment stage, but would not be able to choose the action differentiating these policies in the dynamic game.

The power consistency condition has several interpretations discussed at length in the paper. For example, if an important decision requires unanimity in the dynamic setting, the simple majority rule should not be used at the commitment stage to compare policies differing only with respect to this decision. Here, power consistency reflects the notion that the importance of the decision is the same whether it is considered in the dynamic game or at the commitment stage. In other settings, power consistency captures a notion of fairness toward future generations. The condition prohibits current society members from committing to future actions which are contrary to the interest of future society members, who would normally be the ones deciding on these actions. Power consistency may also capture a notion of liberalism similar to the one described by Sen (1970): the social ranking of policies should respect the preferences of individuals who would naturally be making decisions in the dynamic setting. Violations of power consistency are also natural in some contexts. In particular, it is well-known that commitment is valuable for a time-inconsistent agent. We will explain how time inconsistency creates a particular form of power inconsistency which advantages the first-period self (or preference) of the agent.

The observation that commitments can lead to indeterminacy has been made before. In particular, Boylan and McKelvey (1995), Boylan et al. (1996) and Jackson and Yariv (2014)

[4]Thanks to this assumption, the cycles which may arise among state-contingent policies have nothing to do with possible cycles in any given period. They also result in equilibrium uniqueness which simplifies the statements of the paper.

show, when agents have heterogeneous discount factors, that no agreement can be reached over consumption streams because no Condorcet winner exists in their setting. Like us, these authors are concerned that the absence of a Condorcet winner weakens the applicability and value of commitment. By contrast, Acemoglu et al. (2012) and Acemoglu et al. (2015) provide single-crossing conditions on agents' preferences under which the equilibrium is undominated and a dynamic median voter theorem applies.[5]

Unlike these earlier works, our result does not affirm or negate the existence of a Condorcet winner among policies. Rather, it provides a necessary and sufficient condition – power consistency – under which the institution of commitment fails to resolve equilibrium inefficiency, either because the equilibrium was a good policy to begin with, or because commitments lead to indeterminacy. This result may play a cautionary role in settings where committing to *some* policy could improve upon the equilibrium. Kydland and Prescott (1977) already emphasized the value of commitment in macroeconomic settings and the observation that political equilibria are dominated by some specific commitments has been frequently emphasized until now (e.g., Strulovici (2010) and Dziuda and Loeper (2015)). These observations characterize equilibrium inefficiency in specific dynamic settings. However, our result shows that the option to commit is not necessarily a cure for inefficiency. In fact, unless there is a reason why the switch from dynamic decisions to comparing policies puts a different group in charge, or certain policies are a priori ruled out, commitment cannot help.

Of particular relevance to our result, the literature on agenda setting has pointed out long ago (e.g., Miller (1977)) that if the winner of a sequence of binary majority votes over alternatives depends on the order in which alternatives are compared, then there is no Condorcet winner among these alternatives.[6] Our setting is different because the binary decisions of our dynamic game are not the same as the elimination of alternatives in the agenda setting literature. When there is no uncertainty, any policy reduces to a single path in the dynamic game and can be identified with its unique terminal node. In this particular case, policies can then be identified with the "alternatives" of the agenda setting literature. With uncertainty, however, this relation breaks down because policies are state-contingent plans which can no longer be identified with terminal nodes. Choosing among policies at the commitment stage is thus no longer equivalent to making a sequence of binary choices in the dynamic game. For example, one may construct examples (Appendix B) in which reversing the order of moves in the dynamic game does not affect the equilibrium, and yet the equilibrium is Pareto dominated by a commitment to some other policy. We will also explain how our treatment of arbitrary decision rules further distances our results from this literature.

After describing our main result for the simple majority rule in Section 2 and the general

[5] Appendix E.1 provides similar conditions for our setting.

[6] In a static choice problem, Zeckhauser (1969) and subsequently Shepsle (1970) study the existence of Condorcet winners in voting over certain alternatives and lotteries over them. Zeckhauser shows that, if all lotteries over certain alternatives are in the choice set, no Condorcet winner can be found, even if there is such a winner among certain alternatives. In a comment on Zeckhauser, Shepsle demonstrates that a lottery can be a Condorcet winner against certain alternatives that cycle.

case in Section 3, we discuss interpretations of our power consistency condition, and violations thereof, in Section 4. Section 5 presents two applications. The first one concerns search committees and is illustrated by a job market example. It emphasize the role of played by information on our result. The second application concerns reforms and emphasizes the possible role played by appropriate commitment restrictions to improve equilibrium outcomes. Our model focuses on an arbitrary finite horizon. Appendix E.2 shows how our ideas can be adapted to an infinite horizon.

2 Preliminaries: Simple Majority Rule

There are T periods and N (odd) voters. Each period starts with a publicly observed state $\theta_t \in \Theta_t$, which contains all the relevant information about past decisions and events. At each t, a collective decision must be made from some binary set $A(\theta_t) = \{\underline{a}(\theta_t), \bar{a}(\theta_t)\}$. This choice, along with the current state, determines the distribution of the state at the next period. Formally, each Θ_t is associated with a sigma algebra Σ_t to form a measurable space, and θ_{t+1} has a distribution $F_{t+1}(\cdot|a_t, \theta_t) \in \Delta(\Theta_{t+1})$. If, for instance, the state θ_t represents a belief about some unknown state of the world, θ_{t+1} includes any new information accrued between periods t and $t+1$ about the state, which may depend on the action taken in period t. The state θ_t may also include a physical component, such as the current stage of a construction in the Roman metro example.

Let $\Theta = \cup_{t=1}^{T} \Theta_t$ and $A = \cup_{\theta \in \Theta} A(\theta)$ denote the sets of all possible states and actions. Each voter i has a terminal payoff $u_i(\theta_{T+1})$, which depends on all past actions and shocks, as captured by the terminal state θ_{T+1}. A *policy* $C : \Theta \to A$ maps at each period t each state θ_t into an action in $A(\theta_t)$.

If a policy C is followed by the group, then given state θ_t, i's expected payoff seen from period t is

$$V_t^i(C|\theta_t) = E[u_i(\theta_{T+1})|\theta_t, C].$$

Given a policy C and state θ, let C_θ^a denote the policy equal to C everywhere except possibly at state θ, where it prescribes action $a \in A(\theta)$.

Definition 1 (Voting Equilibrium). *A profile* $\{C^i\}_{i=1}^N$ *of voting strategies forms a* Voting Equilibrium in Weakly Undominated Strategies *if and only if*

$$C^i(\theta_t) = \arg\max_{a \in A(\theta_t)} V_t^i(Z_{\theta_t}^a|\theta_t)$$

for all $\theta_t \in \Theta$*, where* Z *is is the policy generated by the voting profile:*

$$Z(\theta_t) = a \in A(\theta_t) \text{ if and only if } |C^i(\theta_t) = a| \geq \frac{N}{2}.$$

The Z is defined by simple majority voting: at each time, society picks the action that garners the most votes. The definition captures the elimination of weakly dominated strategies: at each t, voter i, taking as given the continuation of the collective decision process from period $t+1$ onwards that will result from state θ_{t+1}, votes for the action that maximizes her expected payoff as if he were pivotal.

We assume for simplicity that for each period t, state θ_t, and policy C, each voter has a strict preference for one of the two actions in $A(\theta_t)$. This assumption is standard in the tournament literature.[7]

Because indifference is ruled out and the horizon is finite, this defines a unique voting equilibrium, by backward induction. The proof of this fact is straightforward and omitted.

Proposition 1. *There exists a unique voting equilibrium.*

Commitment and Indeterminacy

Given a pair (Y, Y') of policies, we say that Y dominates Y', written $Y \succ Y'$, if there is a majority of voters for whom $V_1^i(Y|\theta_1) > V_1^i(Y'|\theta_1)$. A Condorcet cycle is a finite list of policies $Y_0, \ldots, Y_K$ such that $Y_k \prec Y_{k+1}$ for all $k < K$, and $Y_K \prec Y_0$. Finally, X is a Condorcet winner if for any Y, either $X \succ Y$ or X and Y induce the same distribution over Θ_{T+1}.

Theorem 1. *Let Z denote the equilibrium policy.*

i) If there exists Y such that $Y \succ Z$, then there is a Condorcet cycle including Y and Z.

ii) If there exists a policy X that is a Condorcet winner among all policies, then X and Z induce the same distribution over Θ_{T+1}.

Remark 1. *If voters' preferences allow ties, Part i) still holds with a weak Condorcet cycle: there is a finite list of policies $Y_0, \ldots, Y_K$ such that $Y_k \preceq Y_{k+1}$ for all $k < K$, and $Y_K \prec Y_0$. Furthermore, Z continues to be a Condorcet winner in the sense that there does not exist another policy Y such that $Z \prec Y$.*

The proof, which is subsumed by the proof of the general case (Theorem 2) and omitted, may be outlined as follows: If Y is different from the equilibrium Z, then it must take actions which the majority opposes on some set of states. One may thus construct a new policy, Y_1, which is preferred to Y and closer to Z, by switching some action of Y. Repeating this step for Y_1 and proceeding by induction, we can then construct a sequence of policies increasing in the majority ranking and get gradually closer to Z until Z is reached. Thus, if Y was dominating Z, we have constructed a cycle. When there are infinitely many states, the argument must be modified to reach Z in a finite number of step. The actual proof works by lumping states together according to periods and winning coalitions in order to get a finite cycle. The proof shows that, starting from any policy distinct from the equilibrium policy, one can construct

[7] Preference relations across alternatives are assumed to be asymmetric. See Laslier (1997). Without this strictness assumption, most of Theorem 1 still applies to "weak" Condorcet winner and cycle. See Remark 1.

a finite chain of polices, increasing in the majority ranking, going from this policy to the equilibrium.

The cycles predicted by Theorem 1, whenever they occur, may be interpreted as follows: If the population were allowed, before the dynamic game, to commit to a policy, it would be unable to reach a clear agreement, as any candidate would be upset by some other proposal. If one were to explicitly model such a commitment stage, the outcome of this stage would be subject to well-known agenda setting and manipulation problems, and the agenda could in fact be chosen so that the last commitment standing in that stage be majority defeated by the equilibrium of the dynamic game.[8]

Theorem 1 distinguishes two cases: when the equilibrium is undominated and when there is no Condorcet winner. These cases can often coexist in the same model, for different parameter values. This was the case in the Roman metro example, where the equilibrium is undominated for $q \in [0, 1/3] \cup [2/3, 1]$ and no Condorcet winner existed for $q \in (1/3, 2/3)$.

A more positive interpretation of Theorem 1 is that, even when the equilibrium policy is majority-dominated by another policy, it must belong to the top cycle of the social preferences based on majority ranking.[9] In the agenda-setting literature, it is well-known that the equilibrium must belong to the Banks set (Laslier (1997)). This need not be the case here, however, due to the presence of uncertainty, because the dynamic game does not give voters enough choice to compare all policies: the decision set is just not rich enough. In particular, with T periods agents make only T comparisons throughout the dynamic game, but policies, being state-contingent plans, are much more numerous when the state is uncertain. As a result, the equilibrium does not *per se* inherit the Banks-set property.

Another way of understanding the difference between the alternatives compared in the agenda-setting literature and the policies compared in our framework is that a state-contingent policy now corresponds to a *probability distribution* over terminal nodes, and in the dynamic voting game agents do not have rich enough choices to express preferences amongst all these distributions. Put in the more formal language of tournaments, the choice process along the dynamic game may not be summarized by a complete algebraic expression for comparing all policies (Laslier (1997)). These differences are substantial and indeed, the method of proof used for establishing our main theorem is quite different and significantly more involved than the one used in deterministic setting to show that the equilibrium is dominated if and only if there is no Condorcet winner among simple alternatives.

[8]One could also incorporate commitment decisions into the dynamic game, with the state θ_t encoding whether a commitment has been chosen before period t (and if so, which one).

[9]Even then, however, the equilibrium policy may be Pareto dominated by another policy, as in the recruiting application described in Section 5.

3 General Voting Rules and Power Consistency

Collective decisions often deviate in essential ways from majority voting. In the Roman metro problem, for example, civil engineers and archeologists have a special say over the feasibility and importance of preserving uncovered ruins. Another natural example concerns constitutional amendments in the United States, which require a supermajority rule. This section shows that our main result still holds for arbitrary decision rules, under a *power consistency* condition whose meaning and relevance are discussed in detail below.

The formal environment is the same as before except for the structure of political power.[10] Given a period t and state θ_t, the "high" action $\bar{a}(\theta_t)$ might, for instance, require a particular quorum or the approval of specific voters (veto power) to win against $\underline{a}(\theta_t)$. The decision rule may also depend on the current state and, through it, on past decisions. In many realistic applications, some voters may be more influential than others because they are regarded as experts on the current issue, or because they have a greater stake in it, or simply because they have acquired more political power over time.

To each state θ_t corresponds a set $\bar{\mathcal{S}}(\theta_t)$ of coalitions which can impose $\bar{a}(\theta_t)$ in the sense that if all individuals in $S \in \bar{\mathcal{S}}(\theta_t)$ support $\bar{a}(\theta_t)$, then $\bar{a}(\theta_t)$ wins against $\underline{a}(\theta_t)$ and is implemented in that period. Likewise, there is a set $\underline{\mathcal{S}}(\theta_t)$ of coalitions which may impose $\underline{a}(\theta_t)$. These sets are related as follows: $\underline{\mathcal{S}}(\theta_t)$ contains all coalitions whose complement does not belong to $\bar{\mathcal{S}}(\theta_t)$, and vice versa. We impose the following condition: for any coalitions $S \subset S'$ and state θ, $S \in \bar{\mathcal{S}}(\theta) \Rightarrow S' \in \bar{\mathcal{S}}(\theta)$. This monotonicity condition implies that it is a dominant strategy for each individual to support their preferred action, for any given state: they can never weaken the power of their preferred coalition by joining it.

A *coalitional strategy* C^i for individual i is, as before, a map from each state θ_t to an action in $A(\theta_t)$. It specifies which action i supports in each state. Given any profile $\mathbf{C} = (C^1, \dots, C^N)$ of coalitional strategies and any state θ, there are two coalitions: those who prefer $\bar{a}(\theta)$ and those who prefer $\underline{a}(\theta)$ and one of them is a *winning coalition*: it can impose its preferred action.[11] Let $a(\mathbf{C}, \theta)$ denote this action.

Given a policy C and state θ_t, i's expected payoff seen from period t, is given by

$$V_t^i(C|\theta_t) = E[u_i(\theta_{T+1})|\theta_t, C].$$

Definition 2 (Coalitional Equilibrium)**.** *A profile* $\{C^i\}_{i=1}^N$ *of coalitional strategies forms a* Coalitional Equilibrium in Weakly Undominated Strategies *if and only if*

$$C^i(\theta_t) = \arg\max_{a \in A(\theta_t)} V_t^i(Z_{\theta_t}^a|\theta_t)$$

[10] The number of voters need not be odd any more. We do maintain the assumption that decisions are binary in each period to avoid the complications arising from coalition formation with more choices and equilibrium multiplicity.

[11] That is, the coalition of individuals preferring $\bar{a}(\theta_t)$ belongs to $\bar{\mathcal{S}}(\theta)$ if and only its complement does not belong to $\underline{\mathcal{S}}(\theta)$.

for all $\theta_t \in \Theta$, where Z is is the policy generated by the profile: $Z(\theta_t) = a(\mathbf{C}, \theta_t)$).

The definition is the same as for a majority voting, except that now the action that wins in each period is the one supported by the strongest coalition. We maintain the assumption of the previous section that each voter has, for any policy and state θ_t, a strict preference for one of the two actions in $A(\theta_t)$. Because indifference is ruled out and the horizon is finite, this defines a unique coalitional equilibrium, by backward induction (the proof is omitted).

Proposition 2. *There exists a unique coalitional equilibrium.*

Commitment and Indeterminacy

Now suppose that society members are given a chance to collectively commit to a policy instead of going through the sequence of choices in the dynamic game. When can they agree on a policy that dominates the equilibrium? We need to specify the structure of power at the commitment stage. Given a pair (Y, Y') of policies, say that S is a winning coalition for Y over Y' if $Y \succ Y'$ whenever all members of S support Y over Y' when the two policies are pitted against each other. A power structure specifies the set of winning coalitions for every pair of alternatives. Given a power structure and a profile of individual preferences over all policies, one can then construct the social preference relation, which describes the pairwise ranking of every two alternatives: $Y \succ Y'$ if and only if there is a winning coalition S for Y over Y' all of whose members prefer Y to Y'. Our assumptions guarantee that the preference relation is complete.[12]

Given the social preference relation $\succ$, say that a policy Y is a *Condorcet winner* if there is no other policy Y' strictly preferred over Y by a winning coalition. A *Condorcet cycle* is defined as in the previous section with the only difference that $\succ$ is used instead of the simple majority preference relation.[13]

Our main result relies on a consistency condition relating the power structures in the dynamic game and at the commitment stage.

Definition 3 (Power Consistency)**.** *Suppose that Y and Y' differ only on a set $\bar{\Theta}_t$ of states corresponding to some given period t and that S is a winning coalition imposing the action prescribed by Y over the one prescribed by Y' for all states in $\bar{\Theta}_t$. Then, S is also a winning coalition at the commitment stage, imposing Y over Y'.*

Although the power structure at the commitment stage must specify the set of winning coalitions for every pair of policies, the power consistency condition is only concerned with a much smaller subset of those pairs, namely the pairs for which the two policies are identical except on a subset of states in a single period.

[12] Although individuals have strict preferences across any two actions in the dynamic game, they will be indifferent between two policies that take exactly the same actions except on a set of states that is reached with zero probability under either policy. We view such policies as identical and say that they "coincide" with each other.

[13] These generalizations of majority-voting concepts to general tournaments is standard. See, e.g., Laslier (1997).

Theorem 2. *Assume power consistency and let Z denote the equilibrium of the coalitional game.*

i) If there exists Y such that $Y \succ Z$, then there is a Condorcet cycle including Y and Z.

ii) If some policy X is a Condorcet winner among all policies, then X and Z must induce the same distribution over Θ_{T+1}.

Theorem 2, whose proof is in the Appendix, implies that if pairwise comparisons of policies are based on the same power structure as the one used in the binary decisions of the dynamic game, allowing commitment will not lead to an unambiguous improvement of the political equilibrium. While some agenda setter may propose a commitment to resolve political inertia, such commitment can be defeated by another commitment proposal, and so on, until eventually returning to the equilibrium political inertia. While one may find some solace in the fact that the equilibrium policy is part of the top cycle among policies, it may of course be Pareto dominated by another policy and one can choose payoffs to make the domination arbitrarily large.

Remark 2. *As with Theorem 1, a modification of Theorem 2 based on weak Condorcet cycles and a weak Condorcet winners holds when agents are allowed to have weak, instead of strict, preferences.*

Power consistency may be illustrated by a modified version of the Roman metro example. Suppose that the decision of preserving ruins – if construction has started and such ruins were indeed found – is made by an archeological body. Also assume that the archeological body's preference is to preserve the ruins. In this case, the general public will vote against undertaking metro construction if most individuals j have an expected payoff $(1-q)m_j + qt_j < 0$ where $m_j \in \mathbb{R}$ is j's benefit from using the metro and $t_j \in \mathbb{R}$ is j's value from ruin restoration (which may be negative if restoration is financed through taxes). One may ask again if Romans could implement an unconditional commitment to complete the metro line. Power consistency implies in this context that the comparison of policy YM (finish no matter what) and YT (build until ruins are discovered) does not lie with the general public but with the archeological body. Still by power consistency, the choice between YT and N is made according to the simple majority rule, so N would defeat YT. Since also YM defeats policy N (uninterrupted metro construction was assumed to dominate the status quo), we get a Condorcet cycle.

The model of this section, by allowing history-dependent power structures, extends the agenda-setting and tournament literatures, which have assumed (see Laslier (1997) for an overview) that the pairwise ranking of "alternatives" was prescribed by a single binary complete, asymmetric relation (tournament), regardless of how or when these alternatives were compared. In dynamic settings such as ours, where each decision affects the balance of power for future decisions, this invariance assumption is typically violated. In the theory of clubs, for instance, an early decision to admit new members dilutes the power of preceding members and, hence, affects the subsequent comparisons of alternatives.

The Necessity of Power Consistency

When power consistency fails, one may find some policies which are unambiguously preferred to the equilibrium. More precisely, we will say that the power structures used in the dynamic and commitment stages are *inconsistent* if there exist policies Y and Y' and a coalition S such that i) Y and Y' are identical except for a subset $\bar{\Theta}_t$ of states of some given period t reached with positive probability under policy Y (and hence Y'), ii) whenever a state $\theta_t \in \bar{\Theta}_t$ is reached in the dynamic game, S is a winning coalition imposing the action prescribed by Y' over the one prescribed by Y, iii) at the commitment stage, S does not belong to the set of winning coalitions imposing Y' over Y.

Theorem 3. *Suppose that the power structures are inconsistent across stages. Then, there exist utility functions $\{u_i(\theta_{T+1})\}_{i\in\{1,\ldots,N\},\theta_{t+1}\in\Theta_{T+1}}$ and a policy X such that the equilibrium Z is strictly dominated by X and X is a Condorcet winner.*

4 Interpreting Power Consistency

When does power consistency hold?

The simplest instance of our setting is when the same set of agents is making decisions at the dynamic and commitment stages, and these agents are time consistent. In this case, power consistency may be interpreted and justified in the following ways.

Expertise: Some decisions (choosing an energy policy, addressing international conflicts, setting monetary policy, etc.) require specific expertise. For these decisions, the power should lie with experts both when these decisions are made in the dynamic game and when comparing policies which differ only with respect to these decisions.
Liberalism: Some decisions primarily concern specific subgroups of the population (e.g., city or state wide decisions, rules governing some associations, etc.), it seems natural to let these groups have a larger say over these decisions both at the dynamic and the commitment stages. This consideration is related to Sen's notion of "liberalism" (Sen (1970)), a link explored further in this section. This consideration may also be applied to minority rights.

Supermajority: Many constituencies require a supermajority rule to make radical changes to their governing statutes. For example, amendments to the United States constitution require two-thirds of votes in Congress, and substantive resolutions by the United Nations Security Council require unanimity. The rules should treat these radical changes consistently whether they are part of a commitment or arise in the dynamic game.

In several policy applications, such as problems with intergenerational transfers of resources, environmental decisions, and international treaties, commitments involve generations which are unborn when the commitments are made. Whether power consistency holds depends on how one treats unborn generations in the social preference relation.

Intergenerational altruism/liberalism: When a decision primarily concerns unborn generations, the social preference concerning policies that differ only with respect to this decision may, normatively, take into account the preferences of these generations – which may depend on the future state – even though they are absent at the time of commitment. Today's generation is then guided by intergenerational altruism when considering commitments.

Departing generations: Inversely, some agents may die or leave the dynamic game following some actions or exogenous shocks. It is then reasonable to ignore them when comparing policies that differ only with respect to decisions arising after they left the game, which is captured by power consistency.

When is power consistency violated?

At the extreme opposite, another view of future generations is to simply ignore them in the social ranking of policies. This approach violates power consistency, and the current generation will generally find commitment valuable in this case.

Myopic/selfish generation: The current generation ignores the welfare and preferences of future generations. Power consistency is then violated, and this is exposed when the preferences of future generations are in conflict with those of the commitment-making generation.

Time inconsistency: Selfish generations capture a broader time inconsistency problem: the preferences of future decision makers are not reflected in today's preferences. The existence of a relationship between inconsistency and the value of commitment should not be surprising if one considers the case of time-inconsistent agents. Time-inconsistent agents violate power consistency because their initial ranking of social alternatives is not representative of their preferences when they make future decisions. One may think of a time-inconsistent agent as a succession of different selves, or agents, each with their specific preferences. At time t, the t-self of the agent is in power; he is the dictator and the unique winning coalition. When considering commitment at time 0, however, only the initial preferences of the agent are used to rank policies, which violates power consistency.

Commitment is deemed valuable in this case, but only because it is assessed from the perspective of the first-period agent. If one were to take the agent's preferences at various points in time, the value of commitment would be subject to the indeterminacy pointed out in Theorem 2.[14]

These observations extend to multiple agents. For example, a set of perfectly identical but time inconsistent agents would obviously face the same issues as a single time-inconsistent agent, regardless of the voting rule adopted in each period. Again, power consistency is violated when one appropriately treats the future selves of those agents as different agents.

A similar source of time inconsistency concerns institutions whose government changes over time, bringing along different preferences. Old governments are replaced by new ones, and old commitments ignore future governments' preferences, violating power consistency. When governments are elected, their time inconsistency may reflect the time inconsistency of the electorate.

Law of the current strongest: Another form of power inconsistency arises when some agents become more politically powerful over time. Their influence on future decisions in the dynamic game extends above and beyond their power at the commitment stage. These power changes may be foreseeable or random, depending on the economic or political fortunes of individuals at time zero. Regardless of the cause, commitment may be valuable as a way to insulate future decisions from the excessive power gained by a small minority. Power consistency is violated because the evolution of individual power is not included in the commitment decision.

Choosing future voting rules

In some applications (Barbera et al. (2001), Barbera and Jackson (2004)), earlier decisions determine the voting rule used for ulterior decisions. More generally, early decisions can affect each agent's voting weight for future decisions. This possibility is allowed by our framework because the state θ_t includes any past decision and determines the set of winning coalitions at time t. Settings where the future allocation of political power is determined by current agents are studied by the theory of clubs (Roberts (2015)) or in mayoral elections (Glaeser and Shleifer (2005)). Barbera et al. (2001) consider voters deciding on immigration policies that would expand their ranks while Barbera and Jackson (2004) study the general problem of voters deciding today on voting rules that will be used in the future.

We now discuss in the context of an example whether power consistency should be expected to hold and what Theorem 2 means when power is endogenous. We start with a two-period model. In period 1, a first generation of voters, assumed for now to be homogenous, chooses the voting rule for period 2, between simple majority and two-third majority. In period 2, the next generation votes on whether to implement a reform. It is assumed that a fraction $x \in [1/2, 2/3)$ of period-2 voters favor the reform. In this case, the period-1 generation can obtain whichever outcome it prefers for period 2, by choosing the voting rule appropriately.

[14] The agent's preferences in the first period may incorporate his future preferences, and this very fact may be the source of the agent's time inconsistency, as in Galperti and Strulovici (2014). However, agent's future preferences do not *directly* affect his ranking of policies at time 1.

Whether power consistency holds is irrelevant, because period 2 voters really have no control over the outcome as they are split in their preferences and bound by the voting rule chosen by their elders. In particular, one may assume that the condition holds so that the conclusions of Theorem 2 apply. Here, the equilibrium is efficient for the first generation and dominates any other policy from their perspective, so we are in the case where a Condorcet winner exists and coincides with the equilibrium, as predicted by the theorem.

Suppose next that there is a third period, and that the voting rule chosen by the first generation must also be used for period-3 decision, with x taking the same value as in period 2. To make the problem interesting, we assume that the first generation wishes to implement the reform in period 2 but not in period 3. In this case, choosing a voting rule in period 1 cannot provide an efficient outcome from the first generation's perspective and committing to a long-term policy clearly increase that generation's utility. Power consistency is violated because the third generation's power to choose the reform in the third period is not reflected in the social comparisons of policies, which is exclusively based on the first generation's preferences.

Finally, suppose that the three generations are in fact made up of the same individuals at different times. There is a fraction x of people who prefer the reform in the second period and the same fraction x of (partially different) people who support it in the third period. Also suppose that the first-period choice, deciding on which voting rule to use in later periods, is made according to the simple majority rule. If in equilibrium the first-period decision is to use the simple majority rule for future periods, the reform is adopted in both periods. If instead the two-third majority rule is chosen in the first period, no reform is adopted in later periods. Suppose that the two-third majority rule is chosen in equilibrium. This means that there is a majority of individuals who dislike the reform in at least one period, so much so that they prefer the status quo to having the reform in both periods, even though there is also a majority (x) of people who, in each period, prefer the reform to be implemented in that period.[15] If we use the simple majority rule when comparing any pair of policies other than the pairs differing only at one period, there is a cycle across policies: a majority of people prefer no reform at all (Z) to both reforms (Y), but a majority prefers reform in period 1 only (X) to Z, and a majority prefers reform in both periods (Y) to X, so that $Y \succ X \succ Z \succ Y$. Power consistency seems reasonable in this setting: whatever decision is made in the dynamic game reflects the preferences of the population at the beginning of the game. The theorem applies and, since the equilibrium is dominated by reform in either period, we get a Condorcet cycle.

[15] For example, suppose that $x = 3/5$ and the $2/5$ who oppose the reform in any given period dislike it much more than they value the reform in the other period. By taking the sets of reform opponents to be completely disjoint across periods, we get $4/5$ of agents against simple majority rule in period 1, as it would lead to reform being implemented in both periods.

Power consistency and liberalism

Sen (1970) has demonstrated that a social ranking rule cannot be both Pareto efficient and satisfy what Sen calls "Minimal Liberalism": for at least two individuals there exists two pairs of alternatives, one for each individual, such that the individual dictates the social ranking between the alternatives in his pair. By linking social preferences to individual decisions in a dynamic game, power consistency can capture Sen's notion of liberalism as a particular case.

Sen's setting concerns a static social choice problem, in which an "alternative" entails a complete description of all decisions in society. When these decisions (collective or individual) can be represented as a dynamic game, Sen's alternatives correspond to the policies studied here and there are natural settings in which power consistency corresponds to liberalism.

To illustrate, consider Sen's main example which concerns two individuals, 1 (a 'pervert') and 2 (a 'prude'), and a book, *Lady Chatterley's Lover*. The prude does not want anyone to read the book but, should the book be read by someone (for simplicity, Sen does not allow both individuals to read the book), she prefers to be the one reading it. The pervert, by contrast, would like someone to read the book, and would also prefer the prude to read it rather than himself (the rationale being that he enjoys the idea of the prude having to read this subversive book). Let x, y, and z respectively denote the following alternatives: 2 reads the book; 1 reads the book; no one reads the book. The situation is captured in the game represented on Figure 2: 1 first decides whether to read the book, then 2 does if 1 elected not to read the book.[16]

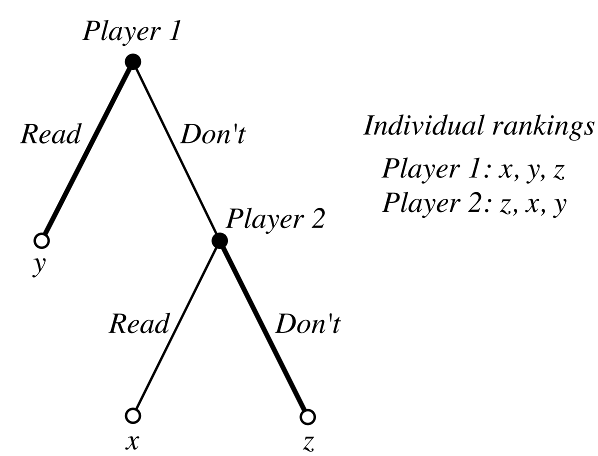

Figure 2: A representation of Sen's game. Individual preferences are indicated from the most preferred to the least preferred alternative.

Power consistency implies that player 2 has the right to choose between reading the book or not. Player 1, too, is entitled to reading the book, regardless of what player 2 does. Thus, power consistency and Sen's version of liberalism are equivalent in this setting. In the

[16] The reverse sequence of moves yields outcome x (the prude player reads the book) and thus does not capture the tension at the heart of Sen's theorem.

coalitional equilibrium of this game, the pervert reads the book and the prude does not (y). Moreover, the Pareto condition of Sen's analysis may be encoded by requiring the unanimity rule for x to win against y. Because x Pareto dominates y given the players' individual preferences, the equilibrium y is defeated by the commitment to a policy in which the pervert does not read the book and the prude does (x). Theorem 2 implies the existence of a Condorcet cycle, which recovers Sen's result on the impossibility of a Paretian liberal.

5 Applications

Search Committees

Consider an economics department deciding whether to fly out a particular job candidate. The faculty initially doesn't know whether the candidate's primary interest lies in macroeconomics or labor economics, but this uncertainty will be resolved during the flyout, should it take place. If the faculty choose not to fly out the candidate, they will settle on a previously seen candidate. Similar models appear in the literature on search committees (Compte and Jehiel (2010), Moldovanu and Shi (2013)). The new candidate is risky because she may polarize the committee. Can commitment help improve the committee's decision?

Let N denote the new candidate and S denote the previously seen, "status quo" candidate. If the flyout takes place, N's field becomes known, and the faculty votes between N and S. The department is divided into three equally sized groups with the following preferences. To a third of the faculty (group *I*), it is important to hire a candidate who will exclusively work on macroeconomics; these members would prefer to make an offer to S over a labor economist. Another third (group *II*) is already convinced about N's value and is willing to choose her over S regardless of her field. The remaining third (group *III*) wishes to choose N over S only if she is a labor economist. Figure 3 represents the payoffs of this game. The utility provided by S is normalized to zero. Adding a twist to the game, we assume that the flyout is intrinsically desirable: even if the department ends up hiring S, it receives a higher payoff from having flown N out, perhaps because getting to know this new faculty and learn about her work is valuable regardless of the hiring decision. Furthermore, there is always a majority ex post who would support making the offer to N.

Despite the benefits of N's flyout, groups *I* and *III* block it for opposite reasons: *I* is concerned that N turns out to be a labor economist, in which case *II* and *III* will collude to make her an offer, while *III* worries that N turns out to be a macroeconomist, in which case *I* and *II* will impose her to the department. Hence, the department fails to fly N out ex ante. This choice is Pareto dominated by the decision to fly N out and then hiring S. In these circumstances, commitment is a tempting solution. For instance, commitment to hire N only if she is a labor economist is majority preferred to the equilibrium. Theorem 1 then implies

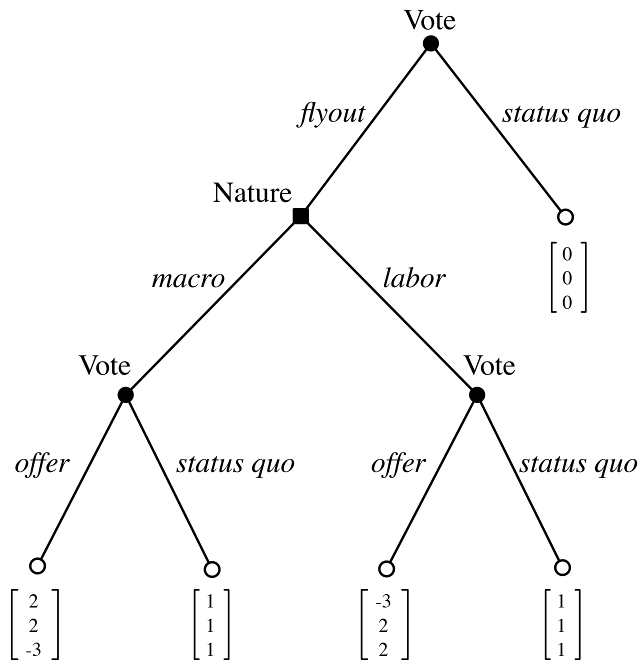

Figure 3: Job Market Game

that there must be a Condorcet cycle.[17]

While decision makers have no clear way out of an inefficient choice, candidate N could fix the problem by revealing her primary interest. By positioning herself clearly as a macroeconomist or a labor economist, she can remove any uncertainty and guarantee herself a flyout. This obversation supports the commonly heard advice that job candidates should avoid mixing fields.

Reforms

A common source of political inertia concerns the avoidance of socially valuable reforms that carry uncertain outcomes and, hence, political risk (Fernandez and Rodrik (1991)). Theorem 1 suggests that the institution of commitment may fail to resolve political inertia. Nonetheless, restricting feasible commitments to anonymous policies provides a natural way of avoiding Condorcet cycles in this application, when agents are symmetric ex ante.

We build on the two-stage setting of Fernandez and Rodrik (1991). In the first stage, citizens of a country decide whether to institute a trade reform. If the reform is undertaken, each individual learns whether he is a winner or loser of the reform. In the second stage, citizens vote on whether to continue the reform, or to implement it if they hadn't done so in the first period. The game is represented on Figure 4. The reform imposes a (sunk) cost c on

[17]One cycle is this: the commitment to fly N out but hiring S in any case defeats not flying N out, but this commitment is beaten by making an offer to N only if she is a labor economist, which is preferred by *II* and *III*. This policy is itself majority-dominated by hiring N in any case, since this yields better outcomes for *I* and *II*. And this last policy is defeated by not flying N out.

each individual that must be borne once, regardless of the duration of the reform. Voters are divided into three groups (*I*, *II*, and *III*) of equal size, and one of the groups is randomly (with uniform probability) chosen as the sole winner from the reform. Individuals in the winning group get a payoff of g per period for the duration of the reform, while remaining individuals lose l per period. If the reform is implemented in the first period and continued in the second, we call it a long-term reform, whereas if it is revoked in the second period, it is a short-term reform.

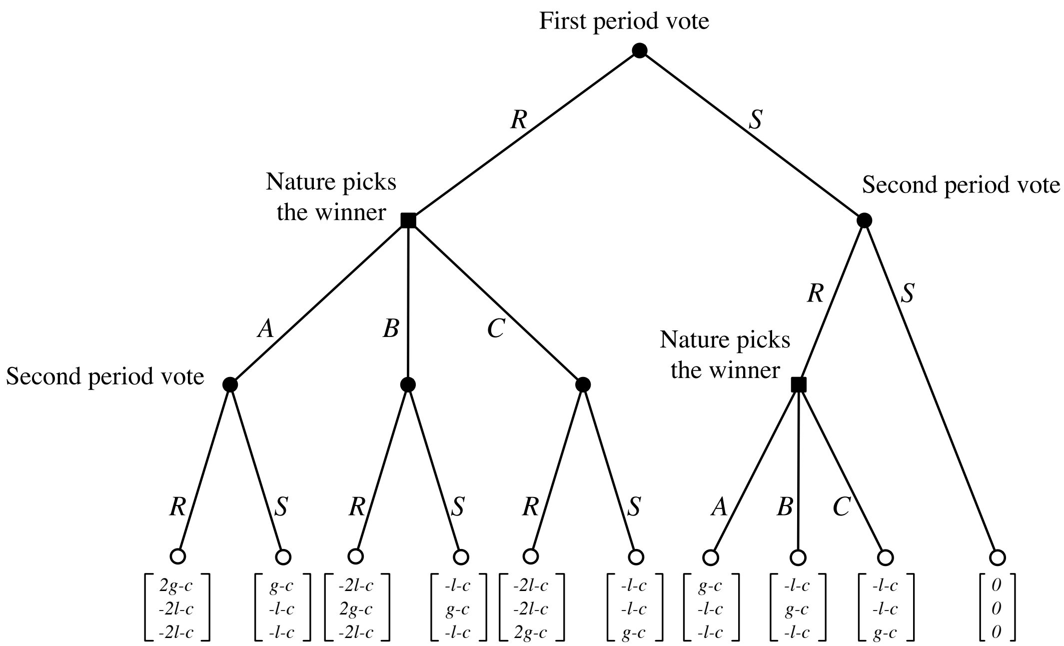

Figure 4: Reform game from Fernandez and Rodrik (1991) for three voters and one winner from the reform.

Provided that g is sufficiently larger than l, the long-term reform is socially valuable. It provides a higher expected payoff to everyone relative to the status quo (not implementing the reform in either period). Commitment to the long-term reform is thus majority preferred to the status quo. However, any initial reform must be revoked in the second period, because two out of three groups find out that they are losers of the reform and have an incentive to end it in the second period. A status quo bias arises if the reform is not implemented at all in equilibrium even though committing to it would be socially beneficial.[18]

Theorem 1 implies that, whenever the status quo bias arises, there must exist a Condorcet cycle over policies, and this cycle involves both the status quo and long-term experimentation. The status quo occurs in equilibrium when $g < 2l+2c$, as in this case the expected payoff from the short-term reform (the reform is ended at the end of the first period, because two groups

[18] "Status quo bias" is defined as the fact that a socially beneficial reform is not implemented despite having a positive expected payoff relative to the status quo. Strulovici (2010) decomposes the notion of status quo bias into "loser trap" and "winner frustration" effects, which respectively mean being stuck with a reform that turns out to be detrimental to oneself, and being unable to implement in the long run a reform which turned out to be profitable to oneself.

of losers are identified) is negative for every type. The expected payoff from the long-term reform is positive for each type, provided that $g > 2l + 1.5c$. Other possible commitments, to short-term reform or to delayed reform in the second period also yield negative expected payoffs. Among these policies, the *long-term reform is a Condorcet winner*, which seems to contradict Theorem 1.

This paradox is explained by the consideration of other possible policies. For example, the commitment to a long-term reform is majority-dominated by the policy which consists of implementing the reform in the first period, and then revoking it only if group *I* is the winner. Groups *II* and *III* strictly prefer this policy to unconditional long-term reform. In turn, this policy is majority dominated by the commitment to continue the reform unless *I* or *II* is the winner. This differs from the previous policy only when *II* is the winner, and in that state of the world *I* and *III* benefit from ending the reform, so their expected payoffs increase. Now that *I* and *II* are at best temporary winners, both prefer the status quo (recall that the expected payoff from short-term reform is negative), which yields a cycle.

These policies, which single out some groups, seem perhaps unfair. In Theorem 1, the Condorcet cycle over policies is constructed over the full set of feasible policies: any plan of action that conditions on states where voting takes place is under consideration. In this application, however, it makes sense to restrict commitment to anonymous policies, and there does exist a Condorcet winner among anonymous policies. This suggests a way of circumventing the negative results of Theorems 1 and 2 by restricting the policy space.

6 Extensions

Random proposers

In well-known agenda-setting protocols, voters may take turns to make collective proposals, and may be chosen deterministically or stochastically to do so. These protocols are compatible with the setting of this paper. For example, for t odd the state θ_t would include, as well as past information, the identity of a proposer who chooses between two collective proposals. At the next, even, period, the new state θ_{t+1} includes the proposal just made and society decides whether to accept the proposal, given the possibility of future proposals.

Non-binary decisions

Our earlier focus on binary decisions in the dynamic game gets rid of Condorcet cycles at the stage-game level, which avoided confusion between these cycles and those, at the heart of our result, which may arise among commitment policies.[19] Many political problems do, in

[19] The approach is also used in the explicit protocol proposed by Acemoglu et al. (2012) (p. 1458). While collective decisions are all binary, that paper allows a player to make proposals among all possible states. This can be easily replicated here by a sequence of at most S periods, where S is the number of states, with each

fact, have this binary structure. For example, choices such as referenda and initiatives take the form of binary decisions. Similarly, lawmakers introduce bills and amendments as "yes" or "no" choices.

With three or more alternatives to choose from at any time, one may attempt to resolve the potential Condorcet cycles by a "binarizing" procedure as in the agenda setting literature. The resulting game then becomes subject to the theorem of this paper: the binary choice sequence leads either to an undominated equilibrium or to an indeterminacy in the ranking of state-contingent policies.

Transfers

Transfers may be explicitly considered in the present setting. For instance, one may include periods at which the binary action corresponds to whether some player i makes a specific transfer to another player j. In this case, i is a 'dictator' over the decision, and the state θ_t keeps track of all past transfers entering players' payoffs at the end of the game.

General utility functions: non separability and past dependence

The utility functions considered in this paper may depend arbitrarily on past states and decisions. For example, they allow decision complementarities across periods and all forms of past-dependence, such as habit formation, addiction, taste for diversity, utility from memories, learning by experimentation, learning by doing, etc. With such features, one must be careful to correctly interpret power consistency.

7 Conclusion

Can political inertia and inefficient equilibria be resolved through the use of commitments? It turns out, unfortunately, that committing to some policy will improve upon the equilibrium policy only if there is a disagreement among decisive factions of society regarding which policy to choose. Allowing commitments thus transforms any inefficiency problem into one of indeterminacy. This finding holds for general state processes and utility functions, allowing social learning, experimentation, and arbitrarily heterogeneous payoffs.

When the power structure underlying collective decisions is extended beyond the simply majority rule, possibly accounting for the time-and state-dependence of those rules, or for changes in the nature of the decisions faced by society, our theorem relies on a power consistency condition which may reflect representativeness, stability, or fairness of the power structure. When power consistency is violated, the value of commitment may be restored,

period corresponding to a new state being presented to the proposer, who makes a "no" decision until being presented by the state that he wants to propose to the group, at which point he votes "yes" and the proposal is made to the group.

often at the expense of decision makers, such as unborn generations, who are poorly or not at all represented at the commitment stage.

One way of circumventing the theorem is to place a priori restrictions on the set of commitments to consider. For example, when the decision problem is *ex ante* symmetric across voters, one may restrict attention to anonymous or symmetric policies, as shown in Section 5. In general, however, voters have heterogeneous preferences from the very beginning of the dynamic problem, which limits the applicability of the anonymity criterion. Other commitment restrictions may help restore the usefulness of commitment institutions and constitutes a natural point of interest for future research.

Appendices

A Computations for the Roman metro example

We solve the game by backward induction, using the elimination of weakly dominated strategies as a refinement. If construction takes place, the ruins are discovered, and a majority consisting of types II and III votes to preserve them at the expense of the metro. Anticipating this, type-I voters initially vote for the project if:

$$E(u_I) = -2q + 1 - q \geq 0 \iff q \leq \frac{1}{3}.$$

II votes for the project regardless of q, because II benefits whether or not an antiquity is found. Type-III voters support the project if:

$$E(u_{III}) = q - 2(1-q) \geq 0 \iff q \geq \frac{2}{3}.$$

Overall, there is a majority in favor of the project at the outset if $q \leq 1/3$ (in which case, it is supported by types I and II) or $q \geq 2/3$ (then, the project is supported by types II and III). But in case $1/3 < q < 2/3$, types I and III join forces, so that a majority opposes the project.

B Limits of agenda setting: ordering-invariant, Pareto-dominated equilibrium

There are two decisions i) U or D and ii) L or R, and three voters. The decisions can be made in any order and are made according to the simple majority rule. The state of the world is revealed between the decisions. The state of the world can take six possible values with probability 1/6 each, and broken down into two components: the role played by each decision (2 possibilities) and, independently, which player is the "sucker" (3 possibilities). We index the states as $\theta_{\omega i}$ where i indicates the sucker and $\omega \in \{A, B\}$ indicates the role of decisions. The payoff structure is as follows:

In state θ_{Ai}, all players get a payoff of 1 if L is chosen regardless of the other decision (U or D). If R is chosen, then all players get a payoff of 2 if D is chosen, while player i gets -100 and other players get 3 if U is chosen. Importantly, these payoffs do *not* depend on the other in which the actions (U/D and L/R) are chosen.

In state θ_{Bi}, the role of actions is reversed as follows: all players get a payoff of 1 if U is chosen, regardless of the other decision (L or R). If D is chosen, then all players get a payoff of 2 if R is chosen, while player i gets -100 and other players get 3 if L is chosen. Again, these payoffs *do not depend* on the other of decisions (U/D and L/R).

The game is constructed in such a way that there is a commitment problem: players would ex ante all prefer to get the payoff of 2, which is always achievable by committing to the action profile (D, R), regardless of the order of these decisions. Without commitment however, there is always a probability 1/2 that in period 2, two players gain by imposing the actions that gives them 3 and gives the other player -100. For example, if L/R is the first decision, then if players choose action R and the action state ω is A, then two players impose U in the second period and the other one gets -100. If instead they choose action L and the action state ω is B, then two players impose action D in the second period. Because of this, players' expected payoff in equilibrium is order the order of $-100/6 \sim -15$.

By symmetry, players get exactly the same expected payoffs if instead the action U/D is made in the first period.

In contrast, committing to (D, R) yields a payoff of 2, regardless of the state of the world. Thus, reordering actions *per se* does not reveal the value of commitment. Of course, as implied by Theorem 1, there is a cycle among commitment policies, as is easily checked.

C Proof of Theorem 2

Fix any policy Y, let $\bar{\Theta}_T$ denote the set $\{\theta_T \in \Theta_T : \quad Z_T(\theta_T) \neq Y_T(\theta_T)\}$. For $\theta_T \in \bar{\Theta}_T$, let $\mathcal{S}_T(\theta_T)$ denote the coalition of individuals who prefer $Z_T(\theta_T)$ to $Y_T(\theta_T)$. Since Z is the coalition equilibrium policy, $\mathcal{S}_T(\theta_T)$ must be a winning coalition given state θ_T. Let $\mathcal{S}_T = \{\mathcal{S}_T(\theta_T) : \theta_T \in \bar{\Theta}_T\}$ denote the set of all such coalitions and p_T denote the (finite) cardinality of $\mathcal{S}_T$. We index coalitions in $\mathcal{S}_T$ arbitrarily from S_1 to S_{p_T}. For each $p \leq p_T$, let Θ_T^p denote the set of $\theta_T \in \bar{\Theta}_T$ for which the coalition of individuals who prefer $Z_T(\theta_T)$ to $Y_T(\theta_T)$ is equal to S_p and for which S_p is a winning coalition. By construction Θ_T^p is nonempty. Consider the sequence $\{Y_T^p\}_{p=1}^{p_T}$ of policies defined iteratively as follows.

- Y_T^1 is equal to Y for all states except on Θ_T^1, where it is equal to Z.
- For each $p \in \{2, \ldots, p_T\}$, Y_T^p is equal to Y_T^{p-1} for all states except on Θ_T^p, where it is equal to Z.

By construction, $Y_T^1 \succeq Y$ because the policies are the same except on a set of states where a winning coalition prefers Z (and, hence, Y_T^1) to Y, and, by power consistency, they can impose Y_T^1 over Y in the commitment stage because for all states in Θ_T^1, S_1 is a winning coalition, it has to be a winning coalition when comparing Y_T^1 to Y. The winning coalition's preference is strict if and only if Θ_T^1 is reached with positive probability under policy Y.

Therefore, either Y and Y_T^1 coincide (i.e., take identical actions with probability 1), or $Y_T^1 \succ Y$. Similarly, $Y_T^p \succeq Y_T^{p-1}$ for all $p \leq p_T$, and $Y_T^p \succ Y_T^{p-1}$ if and only if $Y_T^p \neq Y_T^{p-1}$ with positive probability. This shows that

$$Y_T^{\bar{p}} \succeq \cdots \succeq Y_T^1 \succeq Y,$$

and at least one inequality is strict if and only if the set of states $\bar{\Theta}_T$ is reached with positive probability under Y. By construction, $Y_T^{p_T}$ coincides with Z on Θ_T: $Y_T^{p_T}(\theta_T) = Z(\theta_T)$ for all $\theta_T \in \Theta_T$.

Proceeding by backward induction, we extend this construction to all periods from $t = T - 1$ to $t = 1$. For period t, let $\bar{\Theta}_t = \{\theta_t \in \Theta_t : \quad Z_t(\theta_t) \neq Y_t(\theta_t)\}$. For $\theta_t \in \bar{\Theta}_t$, let $\mathcal{S}_t(\theta_t)$ denote the coalition of individuals who prefer $Z_t(\theta_t)$ to $Y_t(\theta_t)$. Given the continuation policy Z from time $t+1$ onwards, $\mathcal{S}_t(\theta_t)$ is a winning coalition, since Z is the coalitional equilibrium. Also let $\mathcal{S}_t = \{\mathcal{S}_t(\theta_t) : \theta_t \in \bar{\Theta}_t\}$. Letting p_t denote the cardinality of $\mathcal{S}_t$, we index coalitions in $\mathcal{S}_t$ arbitrarily from S_1 to S_{p_t}. Let Θ_t^p denote the set of θ_t's in $\bar{\Theta}_t$ for which the coalition of individuals who prefer $Z_t(\theta_t)$ to $Y_t(\theta_t)$ is equal to S_p and for which this coalition wins. Θ_t^p is nonempty, by construction of S_p. Consider the sequence $\{Y_t^p\}_{p=1}^{p_t}$ of policies defined iteratively as follows, increasing p within each period t, and then decreasing t: for each t,

- For $p = 1$, Y_t^1 is equal to $Y_{t+1}^{p_{t+1}}$ for all states, except on Θ_t^1, where it is equal to Z.

- For each $p \in \{2, \ldots, p_t\}$, Y_T^p is equal to Y_T^{p-1} for all states except on Θ_t^p, where it is equal to Z.

All the constructed policies have Z as their continuation from period $t+1$ onwards. By construction, $Y_t^{p+1} \succeq Y_t^p$ for all t and $p < p_t$ and $Y_t^1 \succeq Y_{t+1}^{p_{t+1}}$ for all t. Moreover, such inequality is strict unless the set of states over which they differ is reached with zero probability.

By construction, the last policy $Y_1^{p_1}$ generated by this algorithm is equal to Z. Let $\{Y_k\}_{k=1}^K$, $K \geq 1$, denote the sequence of *distinct* policies obtained, starting from Y, by the previous construction.[20]

If $Y \neq Z$ with positive probability, then $K \geq 2$. Moreover,

$$Y = Y_1 \prec Y_2 \cdots \prec Y_K = Z.$$

Therefore, we get a voting cycle if $Z \prec Y$, which concludes the proof of part i).

Since Z can never be defeated without creating a cycle, we can characterize a Condorcet winner out of all policies, if it (they) exists, and ii) follows.

D Proof of Theorem 3

We set the terminal payoffs equal to 0 for all policies, except when i) the action sequence until time t has followed policy Y (and hence Y') *and* ii) the state θ_t reached at time t belong to $\bar{\Theta}_t$. In that case, members of coalition S (its complement $\bar{S}$) get 100 (10) if the time t action prescribed by Y' is played and followed by the continuation of policy Y (and hence Y'), with the reverse payoffs if instead the action prescribed by Y is played at time t and followed by the continuation of policy Y. If $\theta_t \notin \bar{\Theta}_t$, everyone's payoff is set to some small $\varepsilon > 0$ for the common continuation of Y and Y', and to zero for all other continuations.[21]

Let Z denote the equilibrium policy. Z must coincide with Y and Y' until period t since it is the only way, for any player, to get a nontrivial payoff. If the state $\theta_t \in \bar{\Theta}_t$, coalition S imposes at time t the action prescribed by Y', so as to achieve its highest possible payoff of 100, and it is in everyone's interest to implement the continuation corresponding to Y' from time $t+1$ onwards, so that even members of $\bar{S}$ get their second highest payoff of 10. Even if $\theta_t \notin \bar{\Theta}_t$, it is also in everyone's interest to follow the common continuation of Y and Y' so as to get ε. This shows that Z coincides to Y'.

We now consider the social comparisons of policies. Clearly, $\bar{S}$ imposes $Y \succ Y'$ since it has the power to do so and this achieves it maximal expected payoff. Moreover, there cannot be any cycle among policies, since Y and Y' Pareto dominate all other policies. Thus, Y is the Condorcet winner[22] among all policies.

[20] We call two policies *distinct* if they induce different distributions over Θ_{T+1}. Policies that differ only at states that are never reached are not distinct.

[21] Although this is largely irrelevant to the gist of the present argument (see Remark 2), one may add arbitrarily small, action-dependent payoffs to break ties at all points of the game for all histories as required by the non-indifference assumption of Theorem 2.

[22] Any other Condorcet winner X must be identical to Y except on a set of histories which has probability 0 when Y (or, equivalently, X) is followed.

References

Acemoglu, D., Egorov, G., and Sonin, K. (2012). Dynamics and stability of constitutions, coalitions, and clubs. *American Economic Review*, 102(4):1446–1476.

Acemoglu, D., Egorov, G., and Sonin, K. (2015). Political economy in a changing world. *Journal of Political Economy, forthcoming.*

Acemoglu, D. and Robinson, J. A. (2005). *Economic origins of dictatorship and democracy.* Cambridge University Press.

Barbera, S. and Jackson, M. (2004). Choosing how to choose: Self-stable majority rules and constitutions. *The Quarterly Journal of Economics*, 119(3):1011–1048.

Barbera, S., Maschler, M., and Shalev, J. (2001). Voting for voters: A model of electoral evolution. *Games and Economic Behavior*, 37(1):40–78.

Battaglini, M. and Coate, S. (2008). A dynamic theory of public spending, taxation, and debt. *American Economic Review*, 98(1):201–36.

Boylan, R. T., Ledyard, J., and McKelvey, R. D. (1996). Political competition in a model of economic growth: Some theoretical results. *Economic Theory*, 7(2):191–205.

Boylan, R. T. and McKelvey, R. D. (1995). Voting over economic plans. *American Economic Review*, 85(4):860–871.

Compte, O. and Jehiel, P. (2010). Bargaining and majority rules: A collective search perspective. *Journal of Political Economy*, 118(2):189–221.

Dziuda, W. and Loeper, A. (2015). Dynamic collective choice with endogenous status quo. *Journal of Political Economy, forthcoming.*

Fernandez, R. and Rodrik, D. (1991). Resistance to reform: Status quo bias in the presence of individual-specific uncertainty. *American Economic Review*, 81(5):1146–55.

Galperti, S. and Strulovici, B. (2014). From anticipations to present bias: A theory of forward-looking preferences. Working paper, Northwestern University.

Glaeser, E. and Shleifer, A. (2005). The Curley effect: The economics of shaping the electorate. *Journal of Law, Economics, and Organization*, 21(1):1–19.

Jackson, M. O. and Yariv, L. (2014). Collective dynamic choice: The necessity of time inconsistency,. *American Economic Journal: Microeconomics, forthcoming.*

Kahn, G. (January 27, 2007). When Rome builds a subway, it trips over archaeologists. *The Wall Street Journal*, Sec. A p. 1.

Kydland, F. and Prescott, E. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–491.

Laslier, J.-F. (1997). *Tournament Solutions and Majority Voting*. Springer, Berlin.

Miller, N. R. (1977). Graph-theoretical approaches to the theory of voting. *American Journal of Political Science*, pages 769–803.

Moldovanu, B. and Shi, X. (2013). Specialization and partisanship in committee search. *Theoretical Economics*, 8(3):751–774.

Roberts, K. (2015). Dynamic voting in clubs. *Research in Economics, forthcoming*.

Sen, A. (1970). The impossibility of a P aretian liberal. *Journal of Political Economy*, 78(1):152–57.

Shepsle, K. A. (1970). A note on Zeckhauser's "majority rule with lotteries on alternatives": The case of the paradox of voting. *Quarterly Journal of Economics*, 84(4):705–09.

Strulovici, B. (2007). Voting and experimentation. Working paper, Oxford University.

Strulovici, B. (2010). Learning while voting: Determinants of collective experimentation. *Econometrica*, 78(3):933–971.

Volokh, E. (2003). The mechanisms of the slippery slope. *Harvard Law Review*, 116(4):1026–1137.

Zeckhauser, R. (1969). Majority rule with lotteries on alternatives. *Quarterly Journal of Economics*, 83(4):696–703.

E Online Appendix

E.1 Existence of a Condorcet Winner

Terminal States

The setting is specialized as follows: there is finite set S of alternatives. The dynamic game starts with some status quo alternative s_0. Thereafter, in each period, a new alternative in S is proposed against the last accepted alternative. If the challenger is accepted, it becomes the new 'status quo.' Otherwise, the previous status quo is pitted against a new challenger. The state θ_t corresponds to the current status quo. So, without loss, $\theta_t \in S$. It is assumed that the proposal protocol is flexible enough so as to guarantee that each possible sequence of S states may arise under some policy as the succession of status quo. This is achieved for instance if the states are ordered from s_1 to s_S and, the number T of periods is such that $T > S^2$, and at each time $t \leq T$ the challenger is the state s_k where k is the remainder in the Euclidean division of t by S.

The payoff of each player i is determined by the state s at the end of the last period. Formally, i's payoff u_i is a map from S to $\mathbb{R}$. Therefore, any policy reduces, in terms of payoffs, to the last alternative. In particular, there is a Condorcet winner among policies if and only if there is a Condorcet winner among states. As it turns out, this exact condition has been studied by Acemoglu et al. (2012), hereafter (AES). In particular, part i) of their Assumption 2 is identical to ruling out Condorcet cycle.[23] Theorem 4 in AES provide sufficient conditions for acyclicity. By a direct translation of their result, we obtain a sufficient condition for the existence of a Condorcet winner. For completeness, we recall the assumptions of AES:

Say that a set $\mathcal{C}$ of coalitions is *regular* if it satisfies the following properties: i) For any $W \in \mathcal{C}$ and $W \subset W'$, we have $W' \in \mathcal{C}$, ii) For any $W, W' \in \mathcal{C}$, $W \cap W' \neq \emptyset$.

For the sufficient condition, we order the set of players from 1 to N. The preference profile $\{u_i(\cdot)\}_{i\in N}$ obeys the *single-crossing property* if for any ordered states $s < s'$ and ordered players $i < j$, $u_i(s') \geq (>) u_i(s)$ implies $u_j(s') \geq (>) u_j(s)$.

Proposition 3. *Suppose the players satisfy the single-crossing property and that for each state s, the set of winning coalitions is regular. Then, the equilibrium of the coalitional game selects the unique Condorcet winner among the S alternatives.*

Proof. The assumptions of Theorem 4 in AES are satisfied.[24] Theorem 4 then implies acyclicity of coalitional preferences which in turn guarantees the existence of a Condorcet winner s^* among alternatives and, by equivalence, among plans. Theorem 2 then implies that the equilibrium selects that policy and hence s^* as the terminal outcome. □

AES also provide conditions – single-peaked preferences and a nonempty intersection condition concerning winning coalitions across different states – under which the stable set is Pareto efficient (Theorem 3). The conditions may readily be adapted to our setting to guarantee that the equilibrium yields a Pareto efficient outcome (regardless of whether it is a Condorcet winner).

Single-crossing property for policies

Interpret the succession of binary decisions of the dynamic game as society's gradual positioning on a single issue, captured on a one-dimensional spectrum, with $\underline{a}_t$ (resp., $\bar{a}_t$) being a move to the left (right). A state-contingent policy thus corresponds to the gradual refinement of the collective position. The policy $\underline{a}_1, \underline{a}_2$ in a two-period problem is clearly more "left" than the policy $\bar{a}_1, \bar{a}_2$. The comparison of $\underline{a}_1, \bar{a}_2$ with $\bar{a}_1, \underline{a}_2$ depends on the relative magnitude of first and second move. In some settings there is a clear ranking of these magnitudes. For example, the first decision may be of first-order importance, while the second decision concerns a finer point. More generally, policies may be ranked lexicographically: we have $Y \prec Y'$ if the first

[23] In both settings, the set of winning coalitions only depend on the current state.

[24] The assumption, maintained throughout the present paper, that players are never indifferent across actions implies that Assumption 6 in AES is satisfied.

period t for which there exists a state θ_t such that $Y_t \neq Y'_t$, we have $Y_t = \underline{a}_t$ and $Y'_t = \bar{a}_t$ for all the states of period t on which Y and Y' are distinct. This ordering captures the idea that earlier decisions matter more for the policy than later refinements. Policies may then be ranked unequivocally from left to right, and the single-crossing property may be imposed over preferences. With the simple majority rule, the median voter theorem guarantees the existence of a Condorcet winner.

E.2 Voting cycles over infinite-horizon policies: An illustration with collective experimentation

The argument used to prove Theorem 1 may be adapted to Markovian settings, applying backward induction to the underlying state rather than to time. To illustrate this, we revisit the model of Strulovici (2010), focusing on three ex ante symmetric voters. Voters choose at each instant between a risky action (the "reform") and a safe action (the "status quo"). As long as the reform is implemented, each voter may receive some good news, which reveals that the reform benefits him; he becomes a "winner" of the reform. Voters who receives no news, called "unsure voters", become more pessimistic about the reform and wish to abandon it in favor of the status quo. Reform is then abandoned if reform winners did not gain the majority. If the safe action is chosen for some belief of unsure voters, no new learning occurs and experimentation is forever abandoned.

The utilitarian policy, denoted by Y, is ex ante strictly preferred to the equilibrium policy, Z (Theorem 6 in Strulovici (2010)).[25] The equilibrium policy is characterized by cutoffs $p(0) > p(1)$ such that experimentation stops when unsure voters' belief p drops to $p(k)$ and no more than k winners have occurred by then (Theorem 1). The utilitarian policy is determined by similar thresholds $q(k)$ (Theorem 2). From Theorem 3, $q(k) < p(k)$ for $k = 0, 1$. Intuitively, experimentation is more valuable to the social planner than to unsure voters, because he includes the utility of winners in his welfare function and has a higher option value of experimentation than individual voters who have to share power. Suppose that $q(1) = 0$ – it is efficient to play the risky action forever from the moment that one winner has been observed. This condition is equivalent (Theorem 2) to $g \geq 3s$, where $g > 0$ is a winner's flow payoff from the reform and $s \in (0, g)$ is everyone's flow payoff with the status quo. This parametric condition is imposed hereafter.

The thresholds $p(1)$ and $q(0)$ solve the following equations:[26]

$$p(1) = \frac{\mu s}{\mu g + (g - s) + p(1)g - s} \tag{1}$$

$$q(0) = \frac{\mu s}{\mu g + (g - s) + 2(q(0)g - s)}. \tag{2}$$

Because $p(1)$ and $q(0)$ are strictly below the myopic cut-off s/g (Theorem 1), this implies that $q(0) > p(1)$.[27] This implies $p(0) > q(0) > p(1)$.

We now construct a cycle, based on the following modifications of the utilitarian policy. Let Y_1 denote the policy that coincides with the social optimum, Y, except that experimentation stops at the threshold $p(1)$ if the only winner observed by that time is Voter 1. Voters 2 and 3 prefer this policy to the utilitarian policy: conditional on Voter 1 being the only winner by the time $q(0)$ is reached. In fact, this policy maximizes Voters 2 and 3's common expected utility. To see this, notice that the utilitarian cutoff $\tilde{q}$ for Voters 2 and 3, starting from $q(0)$, and given that any winner will cause L to be played forever, is characterized by the indifference equation[28]

$$\tilde{q}g + 2\tilde{q}\lambda \left(\frac{1}{2} \left(\frac{g}{r} + \tilde{q}\frac{g}{r} \right) - \frac{s}{r} \right) = s.$$

[25] We refer the reader to the original paper for the results and notation used in this section. Theorem numbers also refer to the original paper

[26] See equation (17), p. 964, for $N = 3$ and $k_N = 1$ and equation (7), p. 947, for $N = 3$, $k = 0$, and $W(1, p) = \frac{1}{3}(\frac{g}{r} + 2p\frac{g}{r})$.

[27] If $p(1) \geq q(0)$, then we would have $p(1)g - s \geq q(0)g - s > 2(q(0)g - s)$, where the strict inequality comes from $q(0) < s/g$. This would imply that the RHS of (1) is strictly smaller than the RHS of (2) and, hence, that $p(1) < q(0)$.

[28] See Equation (7) on p. 947 with $N = 2$ and $k = 0$, and $W(1, p) = 1/2(g/r + pg/r)$.

The first term on the LHS is the average flow payoff for Voters 2 and 3 when the risky action is played. The second term is their average utility jump if one of them becomes a winner, multiplied by the flow probability of that event. The RHS is their flow payoff with the safe action. The last equation is identical to (1) and thus yields the same cut-off.[29] Therefore, $Y_1 \succ Y$.

Consider now the policy Y_2 that is identical to Y_1 except that experimentation stops at $p(1)$ if only Voter 2 has become a winner by that time. This policy is preferred by Voters 1 and 3 over Y_1, by the same reasoning, so $Y_2 \succ Y_1$. Finally, Y_3 is obtained by modifying Y_2 so that experimentation stops at $p(1)$ if only Voter 3 is a winner by that time. Again, $Y_3 \succ Y_2$.

By construction, Y_3 stops experimentation if a single winner has been observed by the time $p(1)$ is reached, as does the equilibrium policy Z. The only difference between Y_3 and Z is that Z stops experimentation earlier, at the threshold $p(0) > q(0)$, if no winner has been observed, whereas Y_3 stops experimentation at $q(0)$. However, subject to the constraint that experimentation stops when p reaches $p(1)$ and at most one winner occurred, and continues forever otherwise, the socially optimal policy is actually to stop at $p(0)$, not $q(0)$. To see this notice that, due to the constraint, the continuation value of a voter who becomes a winner is precisely the continuation value w under the equilibrium policy, and the continuation value for unsure voters when another voter becomes a winner is the equilibrium continuation value, u.

Formally, the socially optimal cut-off $\hat{q}$ for stopping experimentation when no winner has been observed, given the constraint on continuation play, solves the indifference equation

$$3\hat{q} + 3\lambda\hat{q}\left(w(1,\hat{q}) + 2u(1,\hat{q}) - 3\frac{s}{r}\right) = 3s.$$

The first term of the LHS is the aggregate expected flow payoff for the three voters if L is played, and the second term is their aggregate jump in utility if one of them becomes a winner, multiplied by the flow probability of this happening. The right-hand side is the aggregate flow payoff if the safe action is played instead. Dividing by 3, one obtains exactly the characterization of the cut-off $p(0)$.[30] This shows that $Z \succ Y_3$ and yields the voting cycle $Z \prec Y \prec Y_1 \prec Y_2 \prec Y_3 \prec Z$.

The logic of the Roman metro example can also be extended to an infinite-horizon model. In particular, Strulovici (2007) studies experimentation between several restaurants, including a Singaporean restaurant which, in addition to serving Indian cuisine, can teach "voters" about Chinese cuisine.

To reinterpret this example in light of the Roman metro problem, consider the following setting, with three alternatives. S: do not start construction; M: build the metro, even if some ruins are discovered; P: start construction and preserve ruins if some are discovered (do not finish construction). There are also three voters with the following preferences: Voter 1 cares only about the metro; Voter 2 prefers restoring the ruins to building the metro, and building the metro to not digging at all; Voter 3 does not care about the metro, and prefers restoring the ruins if some are found.

Flow payoffs are as follows. The payoff for not starting the project is normalized to 1 for all voters. If the metro is built, Voters 1 and 2 get a flow payoff of 2, while Voter 3 gets 0. If the ruins are restored, Voter 1 gets a flow payoff of -9, Voter 2 gets 2.1 and Voter 3 gets 1.1. Let p denote the ex ante probability of finding some ruins. The analysis of Strulovici (2007) implies the following result.

Proposition When option P is unavailable, the only Majority Voting Equilibrium (MVE) consists in implementing M. If P is added to the set of feasible alternatives and p lies in $(0.1, 0.8)$, the only MVE consists in implementing S.

Proof. In the example analyzed by Strulovici, Voter 3 gets a payoff $\tilde{g} = 0.1$ under M if some ruins are discovered but the metro is built nonetheless (action M). The analysis of that example is identical for all $\tilde{g} \in [0, 0.1]$, and in particular for the value $\tilde{g} = 0$ used in the Roman Metro example.[31] □

[29] Equation (1) has a unique positive solution (the left-hand side is increasing in p which the right-hand side is decreasing p).

[30] See Equation (20), p. 965, with $N = 3$ and $k_N = 1$, and replacing the factor $N - k_N - 2$ by $N - k_N$).

[31] A possible interpretation of a positive $\tilde{g}$, in the context of the Roman Metro, is that Voter 3 gets a small positive payoff in case ruins are discovered, even the metro ends up being built. The paradox applies to that case too.